

Robotic Sensitive-Site Assessment

Sean Winkler, Michael T. Boulet, and Jonathan D. Pitts

Distribution A: Public Release
This work is sponsored by the Defense Threat Reduction Agency Joint Science and Technology Office under Air Force contract FA8702-15-D-001. Opinions, interpretations, conclusions, and recommendations are those of the author and are not necessarily endorsed by the United States Government.

Reconnaissance of sites suspected to contain chemical, biological, radiological, and nuclear threats is currently a manned mission—one that exposes humans to health risks and that relies on imperfect human observations. Lincoln Laboratory participated in an advanced technology demonstration program aimed at improving the safety and efficacy of this manned mission by supplementing it with robotics-based reconnaissance. The resulting prototype system consists of an unmanned ground vehicle equipped with an integrated sensor suite that relays mission-critical data to a web-based user interface viewable by geographically distributed stakeholders.

» Teleoperated unmanned ground vehicles (UGV) were first employed by the Department of Defense (DoD) for explosive ordnance disposal following the 1992–1995 Bosnian War. Since that time, UGVs have become valuable warfighting tools in hazardous operational settings that pose significant risk to the health and lives of U.S. military personnel. Between 2000 and 2010, the DoD began to explore the utility of UGVs in counter–chemical, biological, radiological, and nuclear (CBRN) operations. One such effort, the CBRN Unmanned Ground Reconnaissance Advanced Concept Technology Demonstration, resulted in the successful integration of radiological and chemical sensors on UGV platforms [1].¹ Although these point sensors represented the state of the art at the time, they were only able to detect radiological contamination and moderate- to high-vapor-pressure chemical agents, the latter by sampling contaminated ambient air.

As the CBRN threat evolved, the DoD continued to develop sensing technologies and countermeasures to defend against CBRN agents. In particular, the DoD needed to improve the detection of low-vapor-pressure, persistent (i.e., remains intact on surfaces for a long time rather than evaporating) chemical agents (e.g., nerve agent VX), which present a weak or nonexistent signature to traditional chemical-vapor point sensors. This need led to a new generation of portable chemical detectors based on Raman backscatter technology. In contrast to point sensors, Raman spectroscopy–based detectors operate in standoff mode, directly interrogating a contaminated surface with a laser and measuring the spectrum of the backscattered energy. Low-vapor-pressure chemical threats can be identified in this manner by comparing the received spectrum against a threat signature library on board the sensor. To investigate the operational utility of these next-generation standoff detectors, the Rapid Area Sensitive-Site Reconnaissance (RASR) Advanced Technology Demonstration was initiated by the Defense Threat Reduction Agency Joint Science and Technology Office.

Sensitive-Site Assessment Mission (*heading level 1*)

The RASR program was intended to improve the efficacy and safety of the manned sensitive-site assessment (SSA) mission by combining UGV technology with standoff sensors capable of rapidly detecting and identifying persistent chemical threats. A sensitive site is any location with special economic, intelligence, diplomatic, or military significance. The overall objective of the SSA mission is to conduct reconnaissance to support an actionable decision regarding future exploitation, surveillance, destruction, or abandonment of a site. Mission tasks include identifying hazards, determining the site’s purpose, and characterizing the physical environment of the site with maps and photographs. The RASR

¹ The iRobot 510 PackBot was outfitted with the HazMat Detection Kit, an outgrowth of the CBRN Unmanned Ground Reconnaissance effort, to detect radiological contamination following the Fukushima Daiichi nuclear plant disaster in 2011 [2].

Advanced Technology Demonstration focused on indoor sites (e.g., manufacturing facilities, chemical storage locations, illicit drug laboratories) suspected to contain hazardous chemicals, including chemical warfare agents and toxic industrial chemicals.

Chemically contaminated sites present special challenges to the SSA mission because of the potential for severe injury to or death of mission personnel and the difficulty in isolating, classifying, and identifying substances encountered in such environments. One of the first tasks undertaken by Lincoln Laboratory in support of the RASR program was a mission analysis to better understand the concept of operations (CONOPS) for manned SSA. This analysis could then inform how sensor-equipped UGVs could be incorporated into the SSA mission.

Mission Analysis (heading level 1)

The Lincoln Laboratory team observed DoD CBRN training events conducted by the Marine Air-Ground Task Force and the Chemical Biological Incident Response Force to understand the current tactics, techniques, and procedures associated with manned site assessment. The team also consulted with the regional National Guard Weapons of Mass Destruction Civil Support Teams to further refine their understanding of the mission.

Mission Timeline (heading level 2)

Although specific SSA activities will vary from site to site, a general mission timeline for sites suspected of chemical contamination is depicted in Figure 1. Sites are most typically identified as suspect by ground forces moving through an area. These forces are not trained in handling hazardous materials, so the site is only marked for further assessment. A specially trained SSA unit arrives at the site hours to weeks after a suspected location is identified, depending on the level of concern or the operational priority of clearing the area. Before entering the site, the unit determines the physical boundaries of the hazardous or “hot” zone, establishes an incident command post a safe distance upwind, prepares a decontamination corridor, undergoes a baseline medical evaluation (i.e., to check vital signs), puts on personal protective equipment (PPE), and plans how to approach and access the site. An entry team of three to four members carrying sensors makes the initial entry into the hot zone. While downrange, the team systematically explores the site and relays pertinent observations (e.g., building layout, sensor readings, chemicals found) back to the command post via a handheld radio. Receiving personnel typically record these observations and create a floor-plan sketch on a dry-erase whiteboard (Figure 2). Because of the constraints imposed by their air supply and PPE, entry team members can only spend 15 to 40 minutes downrange. Upon returning from the hot zone, they proceed through decontamination, remove their protective suits, and then debrief the other members of the SSA unit to review, correct, and append the recorded observations and floor-plan sketch. Personnel from the SSA unit may repeatedly reenter the site to completely characterize it before deciding how to proceed. Depending on what the site assessment reveals, the unit may recommend more extensive sample collection (i.e., sensitive-site exploitation) or declare the site safe and abandon it.

information can be sparse, imprecise, and misunderstood when described vocally and relayed over a radio. If subsequent entry teams are expected to pick up where a prior team left off, they must be fully aware of potential hazards and, to avoid duplicative efforts, of areas already assessed by the prior team. Consequently, the debrief process could take up to a half-hour, meaning that the time spent on knowledge transfer between entry teams could be equal to or greater than the time spent by teams inside the hot zone performing the actual assessments. Furthermore, the transient nature of SSA data products (written observations that get erased from the whiteboard and observations from human recall) makes them challenging to fully exploit, share among stakeholders, and archive.

Unmanned Capability for Site Assessment (heading level 2)

Armed with the knowledge gained through the analysis of CBRN training missions, the Lincoln Laboratory team conceived a CONOPS for supplementing the manned SSA mission with a UGV capability to reduce risk to personnel and accelerate the mission timeline (Figure 3). The CONOPS is described as follows: Shortly after arriving onsite, the SSA unit sends a UGV equipped with a sensor suite downrange. The UGV enters the hot zone and begins generating situational awareness products in the form of floor plans overlaid with spatially registered imagery and sensor readings. Meanwhile, the SSA entry team is engaged in setup activities. On the basis of the findings of the UGV, the incident commander can make more informed decisions on the level of PPE required for the initial manned entry, on focal points for the SSA, and on the need for more specialized tools or additional personnel. Once the SSA team enters the hot zone, the UGV may continue to collect data to support the debriefing between subsequent entries. The incident commander and SSA unit can leverage the robot-generated data to request reachback support (i.e., consultation of offsite subject-matter experts), which may include more refined chemical identification, data interpretation, or direction on additional sample collection.

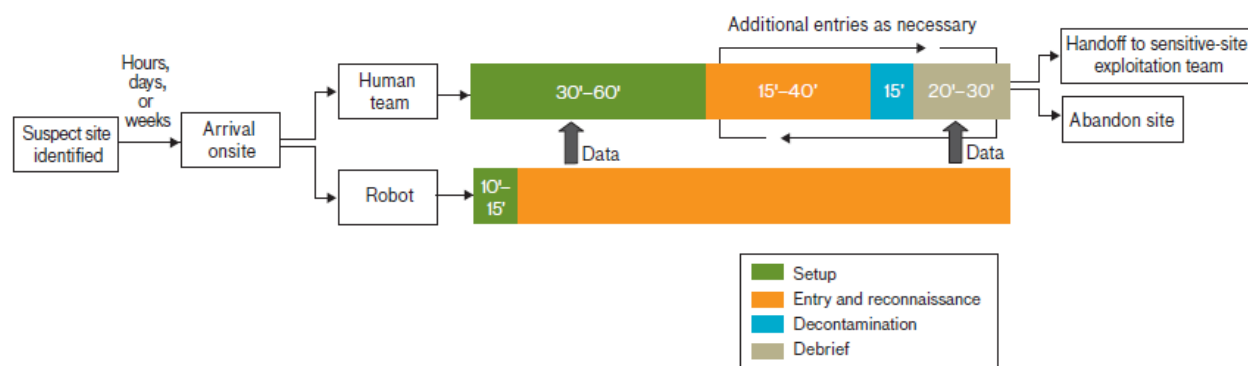


FIGURE 3. A human team and robot operate in parallel to perform reconnaissance of a sensitive site. In contrast to the extensive setup required for manned entry, setup for the unmanned ground vehicle (UGV) is minimal: the UGV has to be unpacked from its case and powered up, and sensors may need to be attached if they are not permanently integrated on the platform. Because of its quick setup, the UGV can begin collecting data in the hot zone soon thereafter the SSA unit arrives onsite. A UGV can provide an early look at the site prior to manned entry and perform persistent reconnaissance when the manned entry team is outside of the hot zone.

Advanced Technology Development (heading level 1)

The industry teams selected by the RASR program focused on integrating their developmental Raman backscatter sensors with UGV platforms, such as the iRobot 510 PackBot and QinetiQ TALON. Additionally, the Defense Threat Reduction Agency asked Lincoln Laboratory to identify and prototype

technologies that could help achieve the envisioned robotic-enhanced SSA operational concept but that were too immature or risky to include within the scope of the industry-led efforts. Such “pathfinder” prototyping is intended to reduce overall program risk by narrowing the gap between academic research concepts and fieldable operational systems. On the basis of lessons learned from the SSA mission analysis, we chose to emphasize the development of novel situational awareness data products enabled by the UGV platform and emerging robotic mapping technology. The primary goals of Lincoln Laboratory’s risk-reduction effort were to provide

- Unmanned (remote) reconnaissance of sensitive sites
- Metrically accurate floor plans combined with immersive panoramic imagery
- Spatially anchored situational awareness and threat sensing
- Human labeling of features in robot-based situational awareness data products (similar to the annotations seen on the whiteboard floor plans)
- Effective sharing and exploitation of data among geographically distributed stakeholders
- Low-latency, interactive technical reachback and collaborative decision making

The remainder of this article is focused on Lincoln Laboratory’s work to achieve these goals.

System Architecture (heading level 2)

A key factor that influenced the RASR system’s architecture design is the distributed and, at times, disconnected set of data sources and stakeholders. The UGVs operate within the hot zone, SSA unit commanders oversee activities at the onsite command post, and stakeholders at higher headquarters (i.e., command centers that provide oversight for a larger military operation) and reachback subject-matter experts may be located thousands of miles from the incident site. In operational settings, network connectivity between these data sources and stakeholders is uncertain and dynamic. The quality of a wireless link between a UGV and command post often varies as the vehicle moves through the environment. Furthermore, the onsite SSA unit may or may not have access to a global information network. A decentralized architecture is needed to support onsite users when access to such a network is unavailable and to support both local and remote users when communication links are available.

As illustrated in Figure 4, the RASR system architecture consists of three components: an enhanced SA payload for unmanned systems, a common operational picture (COP), and a service-oriented architecture (SOA) server. Integrated on the UGV, the situational awareness payload processes and publishes sensor data as the robot explores a site. The COP is a user interface displaying pertinent information in a view shared across all users. Collaboration among stakeholders is supported through various interactive features, including chat, notes, and sketch annotations. The SOA component is the backend infrastructure that receives and stores robot-generated and human-input data and serves these data to several connected COP clients. Multiple SOA instances (i.e., central SOA server, onsite SOA server, and other sites’ SOA servers) provide decentralization for users disconnected from a global network. The system is designed to support multiple concurrent real-time missions and to access archived missions. Details of the payload, COP, and SOA are discussed in the following sections.

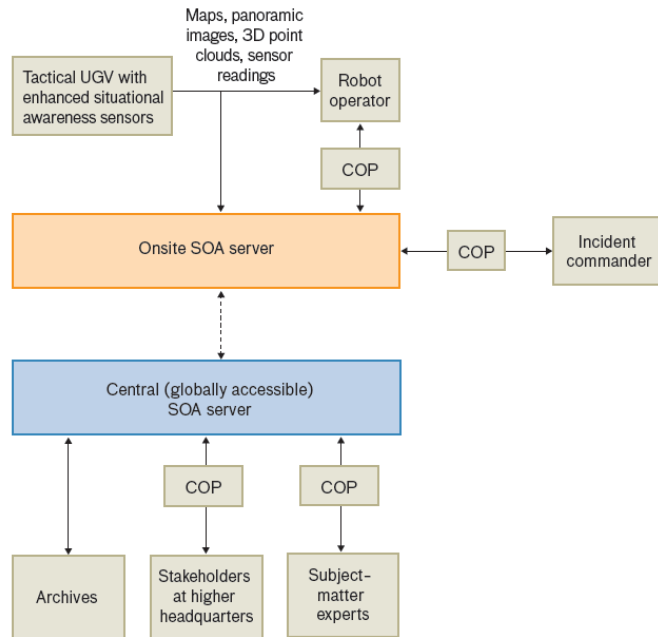


FIGURE 4. The Rapid Area Sensitive-Site Reconnaissance system architecture supports multiple sites and many geographically distributed users with a decentralized design that maintains consistency in information shared between users when interconnecting communication links become available.

Robotic Platform and Payload (heading level 2)

A government-furnished iRobot 510 PackBot tactical UGV was employed as a base mobility platform for the RASR payload. The PackBot is a rugged UGV with tank-style continuous-track propulsion, a seven-degree-of-freedom manipulator arm, and mechanical and electrical interfaces for variable payload integration [4]. With thousands of PackBots deployed in the Iraq and Afghanistan conflicts and many remaining in the current inventory, these UGVs represent the mature ground robot platform capability available today. A PackBot is typically configured for teleoperation over a wireless communications link; using a joystick, a human operator controls the UGV's movement while looking at a video feed from the robot's camera. iRobot was contracted to upgrade the PackBot's internal software to accept commands from the RASR payload over the robot's internal network via a custom message format.

The RASR enhanced situational awareness payload, shown integrated with the PackBot in Figure 5, consists of commercial off-the-shelf sensors, a processor, and auxiliary electronics mounted to a vertical mast rising from the PackBot base. Constructed from T-slotted aluminum framing bars, the payload mast is designed to provide prototyping flexibility. Sensors and other components can be mounted in different places on the mast to evaluate the impact of different configurations on the resulting quality of situational awareness data. Experimentation revealed that a key parameter for sensor placement is height above the floor. Equipment, materials, and other objects relevant to the SSA mission are frequently found resting on tabletops and benches. Mounting sensors on a vertical mast above bench height provides the perspective needed to inspect items of interest. The elevated sensors also improve situational awareness by providing a vantage point more similar to that of a human walking through the environment. Because the mast structure is a single integrated unit that interfaces with the mechanical mounting features and electrical connectors of the PackBot's standard payload, it is easily moved to other PackBot robots. The mast's position over the rear payload bays allows the arm to manipulate objects near the front of the robot.

Onboard arm control software uses a three-dimensional (3D) envelope model that roughly defines the boundaries of the payload to prevent the arm from striking the mast.




FIGURE 5. The enhanced situational awareness payload prototype includes a vertical mast, two lidars, a camera, and a chemical sensor.

Payload Sensor Suite (heading level 3)

The payload sensor suite, powered from the UGV's battery bank, includes a Global Positioning System (GPS) receiver, a hemispherical imaging camera, two scanning laser rangefinders (lidars), and a chemical-vapor sensor. When the system is within view of satellite navigation signals, the GPS receiver estimates the system's global coordinates. These coordinates are subsequently used to geotag the data collection on a map. The hemispherical imaging camera fuses the fields of view of six 0.8-megapixel imagers to create a seamless high-resolution 360-degree panoramic view of the environment. Each lidar estimates the range to objects in the environment (e.g., trees, walls) by emitting an amplitude-modulated continuous-wave laser beam, which reflects off objects in the environment, and by measuring the phase difference of the reflected signal [5]. The lidars scan the laser beam in a plane to produce 1080 distance measurements (with a maximum range of 30 meters) over a 270-degree arc 40 times a second. One of the lidars—the base lidar—is oriented such that its sensing plane is parallel with the ground. This lidar is primarily used for obstacle avoidance and mapping. The other lidar—the mast lidar—has a sensing plane that is perpendicular to the base lidar plane and is mounted to a servo motor. By rotating the mast lidar's sensing plane 180 degrees, the servo causes the plane to sweep through a volume in space to create a 3D representation of the environment. The chemical-vapor sensor continuously samples the air and reports the concentrations of up to five gases, such as volatile organic compounds and carbon monoxide. For this payload prototype, the MultiRAE chemical detector was selected because Raman-based standoff detectors were still under development by the industry teams at the time of the prototyping effort.

Processor (heading level 3)

A laptop functions as the payload's embedded computer, receiving and processing data from payload sensors through standard computer interfaces. The laptop is connected to the PackBot's internal network through which it sends movement commands and accesses the UGV's internal sensors, including camera feeds, arm position, and wheel encoders (i.e., sensors that measure the number of times a wheel turns). The PackBot and payload are typically monitored and controlled by a human operator at the command post. A wireless communication radio links the payload's computer to an operator's console.

Software Architecture (heading level 3)

A key challenge in developing many robotic systems is addressing the size and complexity of the software system needed to read data from sensors, run algorithms, interface with humans, drive actuators, and perform myriad other tasks. Modularity is a useful approach for mitigating software complexity. For example, one module may focus on reading data from sensors while another module may implement a planning algorithm. Compared to a large monolithic application, narrowly focused and well-defined modules are easier to develop, test, and reuse. In order to compose a complete system, however, modules need to be interconnected, sharing information in a consistent way. The RASR payload and operator computer use the Robot Operating System (ROS) software architecture, a widely used open-source robotics-focused library that provides the “plumbing” for connecting executable modules. Additionally, ROS includes a large collection of existing modules implementing commonly used algorithms and drivers for robotic hardware.

The set of executing modules, or nodes in ROS terminology, for the RASR payload is shown in Figure 6. Typically, ROS nodes receive data from hardware or other nodes, perform some processing, and then publish processed data to other nodes or command robot hardware. Data transmitted between nodes are packaged into messages with a well-defined data format and sent over unidirectional streams, or topics. For example, the `/lidar_base` node receives range measurements from the base lidar over the laptop’s universal serial bus (USB) interface, converts those measurements to a standard ROS *LaserScan* message format, and publishes the messages on the `/lidar/base_scan` topic. These laser measurement messages are then received and processed by the `/mapping` node, which, in turn, publishes its results (in this case a floor-plan map), on the `/map` topic. The collection of ROS nodes and topics, which is known as the ROS computational graph, transparently spans multiple computers. The RASR system leverages the multicomputer capability of ROS to execute three components on the operator’s computer: a visualization tool, a joystick hardware interface, and a bridge to the SOA server.

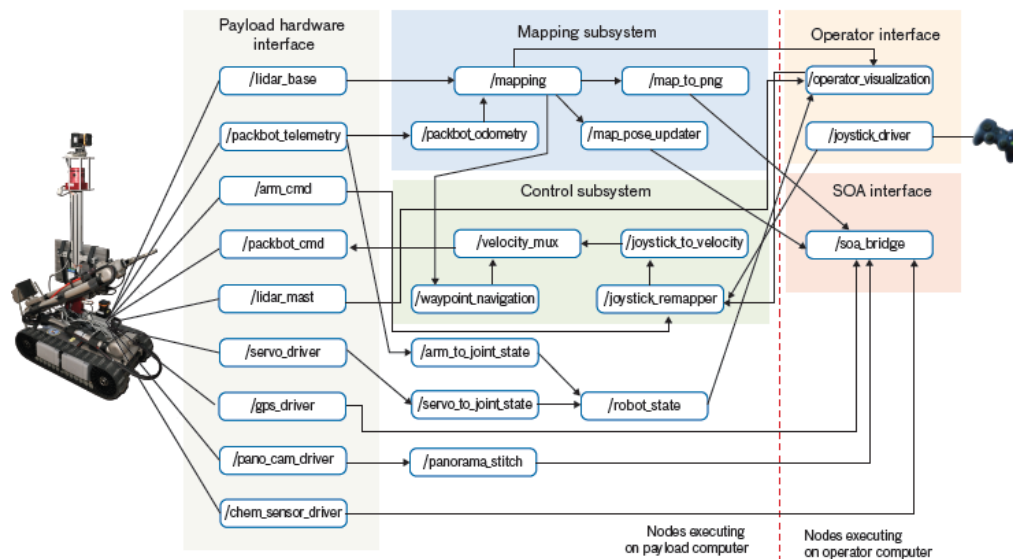


FIGURE 6. The RASR enhanced situational payload computational graph consists of named executing modules, or nodes (rounded boxes), that are connected by data streams, or topics (arrowed lines between boxes). For the sake of clarity, this figure groups related ROS nodes, such as those for the payload hardware interface and the mapping subsystem.

Mapping (heading level 3)

Although all nodes are required to perform remote SSA, the mapping subsystem addresses a key technical

challenge underlying the mission and thus deserves special attention. Maps, such as floor plans or more general world models, are important data structures in robotics. Similar to how humans use road maps in GPS navigation devices, robots use maps to orient themselves in the environment and plan efficient routes to goal positions. For the RASR application, the floor map is also used by human SSA personnel to understand the site's structural layout and to spatially anchor sensor data.

Given a prior map of an environment, such as an as-built floor plan, a robot can estimate its own location and orientation, or pose, by matching features in the map (e.g., doorways or corners) with its sensor observations. However, prior maps are not typically available for sites identified for SSA. On the other hand, if a UGV knows its precise location, creating a map of the environment is a straightforward process of projecting sensor measurements into a suitable spatial representation. However, external navigation signals that might provide a robot with its location (e.g., GPS) are neither sufficiently precise nor generally available in the SSA environment. Therefore, a UGV exploring an unknown site must estimate its own location while simultaneously creating a map. This task, which seemingly presents a causality dilemma (i.e., a chicken-or-the-egg problem), is called simultaneous localization and mapping (SLAM). Solving the SLAM problem has been a focus of robotics researchers for 25 years [6]. The RASR program leverages recent advances in SLAM algorithms to enable robust mapping in real-world, large-scale environments.

Lacking a map, a UGV moving through the environment can estimate its pose relative to its starting location by summing the incremental motion sensed by wheel encoders or an inertial measurement unit over time. Such odometry measurements are noisy because of factors like wheel slip, quantization, and bias drift. Over long robot trajectories, these small errors can accumulate, yielding significant uncertainty in the robot's pose. SLAM algorithms use observations of fixed features, or landmarks, in the environment to correct the accumulated errors. Because no prior map exists, the true locations of the landmarks are not known, and sensors measuring the environment (e.g., scanning laser rangefinder) are also subject to noise. SLAM algorithms use various approaches, such as Kalman and particle filters, to estimate and minimize the global uncertainty in the robot's position and landmark locations.

The SLAM algorithm selected for the RASR effort employs a graph formulation. As shown in Figure 7, the robot's trajectory is sampled at discrete times to create graph nodes (triangles). Odometry measurements and landmark observations (blue and red lines, respectively) create rigid body transformations (i.e., translations and rotations) between poses that form the graph's edges. With perfect sensing, the composition of rigid body transformations between two nodes should be the same regardless of the path through the graph. In practice, with noisy measurements, edges forming two different paths between nodes will conflict. Graph-based SLAM algorithms attempt to reduce the graph's conflicts by adjusting the estimated position and orientation of the nodes. A graph with minimal conflicts represents the best possible map of the environment given odometry and sensor measurement noise.




FIGURE 7. A robot navigating an unknown environment can represent its history as a graph.

New nodes and edges are added to the graph as the robot moves through a site. The SLAM process periodically executes a graph optimization to find a configuration of robot poses that minimizes conflict. Typically, robot pose adjustments are small. Occasionally, observations of a previously seen landmark will shift past robot poses significantly. Such updates present a challenge to the RASR system because situational awareness data are pinned to map coordinates. The */map_pose_updater* module detects large changes in the pose graph and refreshes the location of the data accordingly.

The output of the SLAM process is an occupancy grid map (Figure 8). Similar to an image, an occupancy grid map represents the environment as an array of fixed-sized cells (e.g., 5 cm × 5 cm) with one of three possible states—occupied, unoccupied, or unexplored. Occupied cells correspond to building structures or other obstacles detected by the lidar. Unoccupied cells (consider hallways or open space) are believed to be free of obstructions. The occupancy state of unexplored cells is unknown.




FIGURE 8. A grid map was computed by the Rapid Area Sensitive-Site Reconnaissance mapping subsystem as the robot moved to the right of the figure.

Common Operational Picture (heading level 2)

The COP is a collaboration tool supporting shared situational awareness and decision making for SSA stakeholders. Its design is inspired by human SSA teams' use of dry-erase whiteboards to document observations and to communicate between entry teams. However, the COP extends its reach beyond the local site and current mission. Geographically distributed users may simultaneously access mission data, request and provide analysis of a site, and augment robot-generated data with graphical and textual annotations. An on-scene user, for example, could send images and sensor readings in real time to a laboratory-based subject-matter expert to identify an unknown substance. Or, an analyst at higher headquarters could query archived datasets to see how substances encountered in the field have changed over time.

A screenshot of the COP client interface, which runs in any common web browser, is shown in Figure 9. The interface is divided into three sections: a map-anchored shared data visualization and annotation panel on the left, a workspace in the upper right to record notes and chat with other users, and a panel in the lower right for grouping related data into folders.

The map panel provides spatial navigation tools similar to those of Google Maps. Users can zoom and pan to view details of the occupancy grid map. The interface is touch enabled, so operators with tablets or mobile devices can manipulate the map with familiar touch gestures (e.g., swipe, one-finger tap). Users can augment the map with text and drawing annotations, including icons to indicate doorways or hazards. Locations of sampled payload data, such as chemical-sensor readings or panoramic images, are

represented as colored pins on the map. Selecting a pin reveals additional information (e.g., gas concentration levels for the chemical detector). A layer feature provides a filtering mechanism to control what types of data are included on the map. The map panel also displays 360-degree panoramic images as they are selected and gives users the ability to pan and zoom in a way similar to that of Google Street View (Figure 10).

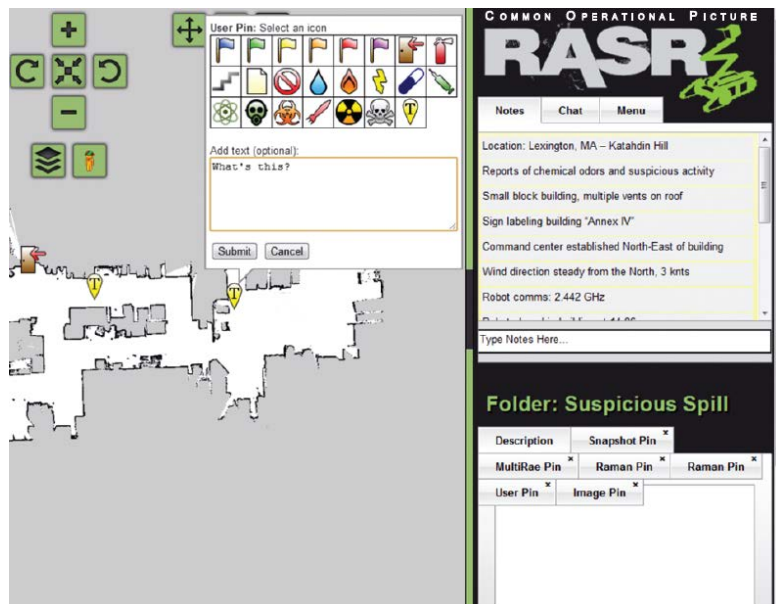


FIGURE 9. The common operational picture interface enables users to place spatially anchored text annotations (yellow markers) on the map panel. A previously placed door icon indicates the location of the building's entryway.

In each mission, the RASR payload collects a large volume of data, potentially resulting in hundreds of pins on the COP interface. Typically, only a subset of sensor measurements is needed for assessing the site's function or its hazards. The folder feature allows users to select, group, and label key measurements. In Figure 10, for example, the user has collected sensor measurements and annotations related to a suspicious spill encountered in a laboratory setting. Each folder is added to the list of layer filters so that a user can limit the displayed pins and annotations to those deemed essential.

The COP can access archived datasets or, if the onsite command post is connected to the global network, display data streamed by a RASR robot payload in real time. Geographically distributed users can observe a mission and inspect sensor data as the robot explores the site, with the map growing as unexplored regions of the environment are revealed by payload sensors. Annotations, folders, and chat messages are also shared in real time with all COP users who have network connectivity.

The user interface builds on standard web technologies to ensure compatibility with most modern web browsers without requiring the installation of plug-ins. The interface uses JavaScript and HTML5 to interact with the user and dynamically update the display as new data are received from the robot payload and other concurrently running COPs. The system extensively leverages open-source JavaScript libraries, including EaselJS and jQuery, to facilitate rapid iteration of the interface layout and functionality on the basis of user feedback.

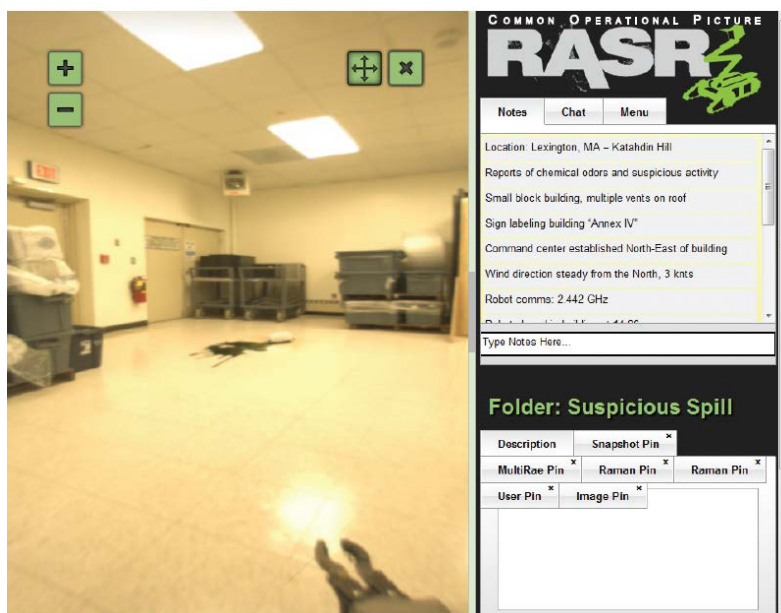


FIGURE 10. The common operational picture displays a 360-degree image captured by the hemispherical camera on the Rapid Area Sensitive-Site Assessment payload. The user can zoom and pan the image to gain a richer understanding of the site. In the lower right panel, the user has created a “Suspicious Spill” folder containing measurements relevant to the spill visible on the floor.

Service-Oriented Architecture Server (heading level 2)

The SOA server provides the backend infrastructure to receive data from robot situational awareness payloads, to archive payload data and user annotations, and to provide data to COP web clients. Several requirements drove the design of the SOA server. The server needed to

1. Be scalable to many concurrent missions and COP users
2. Easily interpret ROS messages produced by the situational awareness payload
3. Support decentralized operation as previously described

A study of common web services and databases led to the selection of Apache’s CouchDB software as the core of the SOA server. In contrast to traditional relational databases, CouchDB enforces few restrictions on the structure of stored data. The RASR system leverages this flexibility to translate ROS messages to CouchDB’s native JavaScript Object Notation document format with little computational cost.

RASR also leverages CouchDB’s replication features to support decentralized operation. Disconnected sites host their own SOA/CouchDB servers, typically on the operator’s computer. The local SOA server receives data from the payload and supports local COP users, such as the incident commander. When connected to the global information grid, the local SOA/CouchDB server transparently synchronizes the recorded payload data with a central SOA server’s data. Remote COP users are then able to access the mission data and add annotations.

Final Demonstration and Impact (heading level 1)

A final demonstration of the end-to-end risk-reduction prototype system was executed in April 2012 at the conclusion of the RASR program. Our robotic platform entered a mock chemical laboratory and performed a site assessment. Data were relayed from the platform to the COP user interface running at the incident command post. Multiple sites across the United States were connected to the COP web interface

so they could view the site assessment in progress. Annotations, interactive data access, chat, and other COP functions were demonstrated in real time.

Future Robotic Sensitive-Site Assessment (*heading level 1*)

While present-day tactical UGVs are very capable on flat ground, their limited agility is a barrier to performing robotic SSA missions in debris-strewn or multilevel environments. The Defense Advanced Research Projects Agency has funded efforts in legged robotics, including the recent Robotics Challenge [7], to improve robot maneuverability in complex environments. However, the transition from wheeled to legged robotic locomotion would pose a challenge to mapping systems like the one used by the RASR payload because these systems commonly assume a level sensing platform and vehicle-like motion. In anticipation of the availability of advanced legged platforms, Lincoln Laboratory teamed with MIT researchers to develop a body-worn mapping system (Figure 11) [8]. A person-portable system serves as an analog for future advanced robotic platforms and could also support collaborative human-robot mapping. The mapping system consists of a vest-mounted sensor suite, similar to that on the RASR payload, and a backpack containing a laptop processor, battery, and barometric pressure sensor. To address the challenges posed by agile robot and human platform motion, the team modified state-of-the-art mapping algorithms to estimate motion by comparing sequential laser range scans, to compensate range returns for gait-induced pitch and roll (rotational movement), and to detect building floor transitions (e.g., ascending a flight of stairs) by measuring ambient air pressure. Demonstrations showed that the system successfully mapped multiple floors of the Ray and Maria Stata Center for Computer, Information, and Intelligence Sciences on MIT campus.




FIGURE 11. In the lightweight body-worn mapping system, sensing and processing components are integrated in a vest-backpack mount. The event push button enables the wearer to “tag” interesting locations on the map.

The mission set of future UGVs may be expanded beyond reconnaissance to include sensitive-site exploitation and other tasks that further reduce the need for personnel to access hazardous environments. For example, with a dexterous manipulation capability, a robot could be tasked with collecting samples and performing analysis using field equipment. Achieving such capability requires significant advances in robotics technology and autonomy algorithms.

Acknowledgments

The RASR Advanced Technology Demonstration is sponsored by the Defense Threat Reduction Agency Joint Science and Technology Office. The authors gratefully acknowledge the technical contributions of Nicholas Armstrong-Crews, Jeanette Baran-Gale, Sarah Curry, Ashley Gavin, Heather Griffin, Anthony Lapadula, Jason LaPenta, Gregory Molnar, and Ray Uttaro to the RASR effort. The man-portable mapping effort is sponsored by Lincoln Laboratory's Advanced Concept Committee. MIT Professors Seth Teller and John Leonard and PhD students Jonathan Brookshire and Hordur Johannsson developed the man-portable mapping technology.

References

1. H.E. Galarraga, P.F. Annunziato, S.M. Funk, and D.E. Green, "Next-Generation Sensor Technology, Now," *Defense AT&L Magazine*, September–October 2009, pp. 12–16.
2. S.M. McPherson, "How Battle-Tested Robots Are Helping Out at Fukushima," *Popular Mechanics*, 18 April 2011, available at <http://www.popularmechanics.com/military/a6656/how-battle-tested-robots-are-helping-out-at-fukushima-5586925/>.
3. H.L. Taylor and J. Orlansky, "The Effects of Wearing Protective Chemical Warfare Combat Clothing on Human Performance," Institute for Defense Analyses, IDA Paper P-2433, August 1991, available at <http://www.dtic.mil/cgi-bin/GetTRDoc?Location=U2&doc=GetTRDoc.pdf&AD=ADA250716>.
4. B.M. Yamauchi, "PackBot: A Versatile Platform for Military Robotics," *Proceedings of SPIE*, vol. 5422, *Unmanned Ground Vehicle Technology IV*, 2004, pp. 228–237.
5. T. Mori, H. Kawata, and S. Yuta, "Light Wave Distance Measuring Apparatus," Hokuyo Automatic Co., Ltd., U.S. patent no. 7,177,014, 13 February 2007.
6. H. Durrant-Whyte and T. Bailey, "Simultaneous Localisation and Mapping (SLAM): Part I, The Essential Algorithms," *IEEE Robotics and Automation Magazine*, vol. 13, no. 2, 2006, pp. 99–110.
7. "DARPA Robotics Challenge," available at www.theroboticschallenge.org/
8. M.F. Fallon, H. Johannsson, J. Brookshire, S. Teller, and J.J. Leonard, "Sensor Fusion for Flexible Human-Portable Building-Scale Mapping," *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 4405–4412.

ABOUT THE AUTHORS



Sean Winkler is a technical staff member in the Intelligence, Surveillance, and Reconnaissance (ISR) Systems and Architectures Group. He joined the Laboratory's Homeland Protection and Air Traffic Control Division in 2001, working first in the Air Traffic Control Systems Group and then the Chemical and Biological Defense Systems Group. He has participated in and led a variety of systems analysis, field measurement, and prototyping efforts in the chemical and biological domain, with particular interest in infrastructure and facility protection. More recently, he has served as a Lincoln Laboratory technical advisor to the Reagan Test Site at the U.S. Army Kwajalein Atoll, working on advanced radar algorithms and sensor fusion. He received a bachelor's degree in mathematics from the University of Texas at Austin and master's and doctoral degrees in applied mathematics from Brown University.



Michael T. Boulet is the assistant leader of the Control Systems Engineering Group, where he focuses on applying autonomous systems technology to address challenging national security problems. He has helped guide the Laboratory's strategic plan for and applied research investments in robotics, autonomy, and human-robot interaction. In his previous role as a technical staff member, he developed perception and control algorithms for a variety of robotics research and prototyping efforts in the ground and air domains. An enthusiast of open-source robotics software, he created the Rapid Robotics and the Rapid Autonomous Complex-Environment Competing Ackermann-Steering Robot (RACECAR) courses for MIT's Independent Activities Period, with the goal of introducing students to advanced autonomous robot capabilities. He joined Lincoln Laboratory in 2001 after graduating from Rensselaer Polytechnic Institute with a bachelor's degree in mechanical engineering. Through the Lincoln Scholars Program, he obtained a master's degree in mechanical engineering from MIT in 2008. In 2012, he received an Early Career Technical Achievement Award for designing a coherent strategy for the Laboratory's autonomous systems effort.



Jonathan D. Pitts is the leader of the Chemical and Biological Defense Systems Group. In this role, he contributes to the management and strategic adaptation of Lincoln Laboratory's portfolio of programs in chemical, biological, radiological, nuclear, and explosives defense. His group specifically focuses on defense architecture development, threat exploitation, gap-filling technology development, and test and evaluation of systems for countering weapons of mass destruction (WMD). After joining the Laboratory in 2003, he pursued the development of environmental monitoring technologies for chemical and biological threats, and operational architectures for such technologies. A major portion of this work concentrated on developing systems that improve U.S. defenses against emerging WMD and transitioning those systems to the field for overseas and homeland use. Prior to joining the Laboratory, he held positions in industry and academia. During that time, he developed optical methods of noninvasive disease diagnosis and tissue engineering. He received a bachelor's degree in chemistry from the University of Rhode Island and a doctoral degree in chemistry from Wesleyan University.